



Integrated Planning and Scheduling for NASA's Deep Space Network – from Forecasting to Real-time

Mark D. Johnston¹ and Jigna Lad²
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Road, Pasadena CA 91109 USA

Over a period of several years, the software systems that plan and schedule the use of NASA's Deep Space Network (DSN) for the projects it serves have been upgraded from a disparate set of decades-old software components, to an integrated suite covering long-range planning and forecasting, all the way to real-time scheduling. The most recent component of this suite is known as LAPS, for Loading Analysis and Planning Software, and is responsible for long-term planning and forecasting, including studies and analysis of new missions, changed mission requirements, downtime, and new or changed antenna capabilities. This paper discusses the architecture of LAPS and its interfaces with other elements of DSN planning and scheduling, its user interfaces, and some lessons learned from development and deployment.

I. Introduction

NASA's Deep Space Network (DSN) is a primary resource for communications and navigation for interplanetary space missions, for both NASA and partner agencies. With three complexes spread roughly evenly around the globe (Fig. 1), the DSN provides services to dozens of active missions. Future growth in mission demand, in both number of missions and science data return, as well as network operations cost drivers, have led the DSN to improve its capabilities for planning and scheduling over all time scales, from multi-year forecasting to real-time rescheduling.

Over a period of several years, the software systems that plan and schedule the use of the DSN for the projects it serves have been upgraded from a disparate set of decades-old software components, to an integrated suite covering long-range planning and forecasting, all the way to real-time scheduling. The planning and scheduling system upgrade has afforded an opportunity to modernize the data model to more accurately reflect the network's current and future capabilities, to improve overall performance and efficiency, and to advance towards better internal and international standardization for data interchange consistent with CCSDS standards as they evolve.

The most recent (2018) addition to the DSN's suite of planning and scheduling software is called LAPS, for Loading Analysis and Planning Software. LAPS is responsible for long-range planning, taking into account factors such as expected mission requirements, antenna capabilities, planned downtime for system upgrades, and the construction of new antennas. LAPS shares a common data model with other parts of the DSN planning and scheduling system, as well as number of interfaces.

This paper describes the architecture of the DSN end-to-end planning and scheduling system, focusing on LAPS and its interfaces with other elements of the system. We describe the new capabilities enabled by LAPS, its user interfaces, and lessons learned from development and deployment. Applicability to communications network planning problems is highlighted, along with usage scenarios such as keeping track of mission requirements over time as a gauge of system demand, and monitoring planned versus actual usage of the DSN.

II. DSN End-to-End Planning, Scheduling and Execution Systems

A. Overview

Fig. 2 shows a block diagram of the DSN planning and scheduling systems, indicating the key mission interfaces and data flows. On the longest time scale, LAPS provides for long-range planning and forecasting, potentially as much as a decade into the future, taking into account anticipated mission requirements and planned DSN asset capabilities.

¹ Principal Scientist, Artificial Intelligence Group, Planning and Execution System Section

² Engineering Applications S/W Engineer, Data Services Group, Planning and Execution Systems Section




Fig. 1. An overview diagram of the DSN antennas at the three complexes in Goldstone, CA, Madrid, Spain, and Canberra, Australia. As of early 2018 there are 12 active antennas, expected to grow to 15 by 2024 with the construction of new antennas and the decommissioning of some older ones.

LAPS users include analysts and mission planners, as well as others involved in long-term planning of DSN activities. At this stage, the missions provide long-term ephemeris information along with their expected utilization of DSN resources.

The Service Scheduling Software (S^3) maintains user-provided detailed scheduling requirements and expands them into specific communications and navigation passes, taking into account resource constraints (antennas and equipment) as well as a wide variety of DSN operational constraints and rules. S^3 supports the scheduling negotiation process, followed by a consensus-based change process with workflows for approval of all schedule modifications by authorized mission representatives. More accurate ephemeris information, along with specific communications and




Fig 2. Schematic of the flow of information through the end-to-end planning, scheduling, and execution monitoring and control systems (top), highlighting interactions with mission users (bottom)

navigation requirements, are provided by the missions during this phase. At the conclusion of this phase, missions receive negotiated DSN allocations that they use to plan and sequence their onboard activities.

The Service Preparation Subsystem (SPS) merges the schedule with mission-provided sequence and ephemeris information to generate detailed DSN ground system sequences of events and a variety of predict calculations to support each scheduled activity. These are distributed to the DSN complexes where they are used by the Network Monitoring and Control subsystem to handle the execution of each pass.

B. Evolution of the DSN Planning and Scheduling Systems

The initial scheduling system development phase, culminating in 2011, led to the creation of the Service Scheduling Software (S³) assembly[1], focused on the mid-range process only — from 2 weeks to about 6 months before execution. Due to the collaborative peer-to-peer negotiation nature of the DSN scheduling process[2], [3], S³ was implemented as a web application with a user interface heavily driven by the need for users from multiple time zones to participate in distributed conflict resolution and change negotiation. S³ is notable in being requirements-driven: project users define their DSN requirements in terms of services, timing, and other constraints and preferences. These requirements are interpreted by the DSN scheduling engine, which generates candidate allocations that satisfy them to the extent possible. The DSN is typically heavily oversubscribed, so there are invariably conflicts that are resolved by human intervention and by users directly negotiating with each other, mediated by the S³ software to manage workflow, change control, and update notifications.

The second development phase extended S³ in 2016 to support real-time DSN scheduling, which enabled a consistent data model from mid-range through real-time. The main enhancement during this development phase was in the area of multiple spacecraft scheduling and operations: for some time, DSN has supported a feature called Multiple Spacecraft Per Antenna (MSPA), in which spacecraft within one beamwidth and at different frequencies can downlink simultaneously, with at most one uplinking at a time. The MSPA capability had been limited to at most two downlinking spacecraft at once. With the 2016 upgrade, the MSPA capability was increased to support up to four simultaneous downlinks and one uplink. While this required changes to more parts of the DSN ground system than just the scheduling software, it was accomplished in a way that was transparent to mission users and has been in successful use since July 2016. The MSPA capability has been heavily (but not exclusively) used for missions at Mars, to minimize antenna utilization while servicing an increasingly large fleet of NASA and international partner spacecraft.

The third development phase, recently completed in 2018, added support for long-range planning and forecasting based on a newly completed component called LAPS – Loading Analysis and Planning Software. Loading analysis and capacity planning has been an essential activity for the DSN, utilizing a succession of software tools[4]–[9]. The LAPS module extended the existing S³ framework to provide an integrated system for all phases of DSN asset planning and scheduling. LAPS provides a web-based user interface for mission phases and events, and for service requirements including timing and repetition. Users can view graphical or tabular renderings of their DSN inputs, in context with other missions' critical events, and with downtime plans over years in the future. Users can save multiple alternative scenarios for parametric trade-off studies. Potential contention intervals can be identified and highlighted, so they can be addressed long before they become critical. Because LAPS is based on the same underlying data model as S³, requirements can be automatically exported from LAPS and used to populate S³. This also makes it straightforward to compare forecasted network utilization with actual timelines as scheduled and executed.

C. Service Scheduling Software

S³ is used to schedule mid-range and near real-time activities, where the boundary between mid- and long-range is not sharp but generally means around 4-6 months from execution. The mid-range process starts when users enter their scheduling requests into the S³ web application, specifying such items as:

- service specifications, in the form of what are called “service aliases”, described further below.
- timing constraints and flexibilities, including track duration, whether extensible or reducible, whether splittable and into how many segments of what minimum duration and separation or overlap, etc.
- visibility constraints such as minimum elevation limits of the spacecraft as seen from DSN antennas, or determined by occultations or shadow events
- relationships to other tracks in terms of minimum and maximum separations
- preferences such as for specific assets, or to schedule early or late in a visibility period

Service aliases are of particular significance in that they specify what combinations of assets are allowable to satisfy a service request. DSN configurations can be complex in that they may be single or multi-antenna (if arrayed), or even multi-complex (e.g. Very Long Baseline Interferometry). There can be requirements for S-, X- and K-band




Fig. 3. A schematic structure diagram of a notional service alias representing a choice between single- and multiple-antenna (array) implementations of the same track. Red highlights the information related to a single-track choice (left branch of “OR”) and blue that related to a two-antenna array choice (right branch). More complex aliases are used to represent up to 4-station arrays, specialized ranging tracks (DDOR), separate uplink and downlink options for multiple spacecraft tracked all at once (MSPA), and maintenance activities that affect an entire complex or the entire DSN at once.

links, and dual band tracking is not uncommon — each requiring a particular complement of equipment, both associated with individual antennas as well as pooled for shared use complex-wide. To capture these configuration options, S^3 introduced the notion of named *service aliases* that describe all possible selectable assets for a service request, as well as certain associated data such as setup and teardown times, default activity descriptions, work category codes, etc.

A diagram illustrating such a service alias is shown in Fig. 3 for a notional Voyager 1 example called “TTC v0”[2]. Here, TTC reflects a Telemetry, Tracking and Command service configuration that can either be satisfied by a single 70-meter track, or by a 2-antenna array of 34-meter antennas at the same complex. Depending on which complex is selected, alternative choices for the specific 70-meter and 34-meter antennas may be made, and then corresponding equipment choices must also be made. Setup and teardown times will also vary. These alternatives will appear very differently in the schedule, as indicated at the bottom of the diagram. The DSN Scheduling Engine component of S^3 is responsible for interpreting service aliases, searching through the tree of possibilities in time, duration, and configuration space, in order to come up with feasible conflict-free schedules. This is a challenging problem and described more fully elsewhere [2].

Service aliases are important in the end-to-end planning and scheduling system because they provide a common mechanism for describing resource requirements and asset configurations that can be shared between long-range planning and the mid-range and near real-time processes. For missions so far in the future that service aliases are not yet defined, it is possible to construct generic aliases reflecting specific DSN service types, for example S-, X- and K-band uplink/downlink/ranging, arrays, DDOR, MSPA, etc.

D. Loading Analysis and Planning Software (LAPS)

Although the primary focus of LAPS is planning over an extended time range to identify periods of asset contention, there are a surprisingly large number of associated functions that are necessarily tightly connected with this forward planning function. Fig. 4 shows a diagram of the LAPS context, indicating key software and user interfaces. These are described in the following.

The LAPS server is a Linux host running NodeJs and MongoDB, with an Apache web frontend and proxy server. The NodeJS web application uses Lightweight Directory Access Protocol (LDAP) authentication and manages access permissions to all the data sets described further below. The MongoDB database server holds two databases used by LAPS: one is dedicated to viewperiod data, i.e. intervals when spacecraft are visible from any of the DSN or partner agency antennas, and the other is used for all other data. Viewperiod data can become very voluminous when covering




Fig. 4. LAPS overview and key functions and user interfaces. See text for details.

long time intervals for extended studies, and the viewperiod database currently holds about 1Gb of compressed visibility data.

Other LAPS functions indicated in the diagram are:

- **DSN Aperture Fee** (costing) calculator and **ULP Generator**: this is an externally accessible web client that allows the user to create, save, and edit planning scenarios for DSN usage, then generate reports of requested antenna loading (in the format of a User Loading Profile, or ULP), as well as a cost profile following DSN rules for computing costs based on antenna usage. This algorithm is quite complex and had in the past been embedded in a spreadsheet that was difficult to use. Such cost information is required to be provided when responding to NASA mission proposal calls, and the ULP is also used in developing DSN service level agreements. The planning scenario file can be imported into LAPS directly, to enable loading studies, and to provide a starting point for missions as they progress through the project lifecycle phases of formulation and implementation.
- **Planning Scenario Editor**: in LAPS, a *planning scenario* is a set of planning events and requests that reflect a mission's anticipated use of the DSN over some time period.
 - Planning *events* can be absolute or relative time intervals (with reference to an event such as an expected launch date), to which requests can be associated. For example, a mission may have standard requirements through a cruise phase, but increased need for coverage around a trajectory correction maneuver. The timing of these periods can be defined as different planning events.
 - Planning *requests* represent patterns of coverage, or customized explicit tracking requests. Attributes of planning requests include a specified service alias, and, optionally, filters that constrain the allowed configurations to a subset of the full suite of possibilities. Other attributes include priority, a description, and selection of one or more tags that can be used to further indicate the intention of the request. Tracking duration can be specified, along with track splitting options (e.g. minimum segment duration and overlap times). Patterns that can be requested include:
 - continuous coverage
 - a daily number of tracks
 - a weekly number of tracks

- tracks that repeat at an arbitrary regular interval
- a collection of tracks specified as relative to the start of a week, that repeat every week
- a track within a specified window

Although a relatively large number of options are available, the intent is to make it easy to translate common request patterns into scenarios with minimum complexity. A request like “during cruise phase, three 8-hour tracks/week of X-band uplink/downlink/ranging” can be quickly specified with a few selections and data value choices.

- A *planning scenario* must refer to a set of assets, downtimes, service users, and various planning parameters. These can either be the baseline sets of these items, or each one can refer to a variant, to support ‘what if’ or trade-off studies.
- **Service User Editor:** the LAPS service user dataset includes information about missions including their names, abbreviations, and descriptions; in addition, service aliases associated with operating missions are recorded here, as well as MSPA capabilities for multiple spacecraft tracking. The LAPS editor allows creating new hypothetical missions for studies and proposals, as well as modifying attributes of existing missions.
- **Asset Editor:** LAPS is cognizant of DSN and partner agency antennas, including 34- and 70-meter antennas at the DSN complexes and similar sized antennas elsewhere in the world. Antennas are grouped into standard subnets (e.g. 34MBWG = 34-meter beam waveguide), used for reporting and resource planning. Antennas have different band capabilities (e.g. X-band uplink and downlink) as well as commission and decommission dates if they are going in or out of service, respectively. These attributes may be edited in LAPS to support studies of different asset scenarios.
- **Downtime and maintenance planning and scheduling:** major downtimes on DSN antennas can last for months and can significantly impact mission users or critical event support. Downtime planning for the DSN is often done years in advance, taking into account user planned activities, engineering personnel and equipment and parts, and cost. LAPS is used to capture new downtime requests, and then used online in downtime planning meetings to review and record decisions. Alternative downtime scenarios can be generated to facilitate impact studies and trade-offs. In addition to extensive downtime editing capabilities, LAPS provides several widely disseminated reports, both graphical and textual:

MAJOR DSN ANTENNA DOWNTIMES




Fig. 5. The LAPS downtime Gantt view shows downtime by complex and antenna at the bottom, along with mission augmented support level activities (“critical events”) in the upper section.

Viewperiods & overlap




Fig. 6. Viewperiod overlap visualizer showing a reference mission viewperiod (top row, and red vertical bands) and the intersection with other missions (other rows). At the right is the calculated percentage overlap. This information is essential in determining and resolving potential contention intervals in the future.

- A downtime Gantt chart view (Fig. 5), showing at-a-glance the downtime scheduled year by year, highlighting downtime by complex or by antenna, and shown simultaneously with critical mission events
- A filterable downtime tabular report, for a selected time frame, specifying antenna, type of downtime, date/time range, and approval status.

Downtime data automatically flows from LAPS to S³ so schedulers always have the latest updates visible as they work on schedule changes. In addition, LAPS exports XML that can be imported to S³ to create engineering and maintenance schedule items, which also helps with ensuring that the latest information is widely available.

- **Viewperiod visualization and overlap calculator:** due to its central role in assessing mission contention, LAPS provides several visualizations of viewperiod alignment and overlap. As seen from the Earth, spacecraft will seasonally tend to overlap in visibility, and at its most extreme, this can make for congestion during those times of day when most missions are in view. LAPS includes a graphic visualizer that shows an arc on a 24-hour clock face that indicates the time of rise and set of each mission at each complex, highlighting the central value. This makes it easy to see when missions are appearing in the same part of the sky, and thus are more congested. This can be viewed on a monthly or weekly basis, as shown in Fig. 7 as part of “DSN This Week”. Additionally, for a specified time range and set of missions, LAPS has a graphical Gantt view (Fig. 6) that shows the precise time of visibility for a set of missions on a set of antennas, with a calculation of the percentage overlap relative to a reference mission.
- **Mission critical events:** A key part of downtime planning is to maintain an accurate and accessible database of mission critical events from various mission phases. This is also used by the DSN complexes for planning their support personnel, and other DSN teams take into account top priority events when planning software deliveries, tests, and shift assignments. LAPS provides a custom mechanism for authorized users to create and track critical events so they can be taken into account elsewhere in the DSN processes. See Fig. 5 for an example of a graphical view of this information; there is also a tabular view with more details about each event.
- **“DSN This Week”:** as an extension of the public “DSN Now” site that shows instantaneous tracking by antenna and mission, LAPS provides an internal 8-day look at the week ahead in a single-page dashboard, based on a correlated visualization of LAPS data combined with the current DSN schedule as published by S³. The “DSN This Week” view includes: mission critical events and corresponding scheduled activity details; viewperiod overlap by mission; tracking time metrics by mission; and planned downtimes at each DSN complex. Each of these has a link to drill down for additional details.
- **Web services:** with its database of planning and forecasting data, downtime, holidays, critical events etc., LAPS provides a valuable interface for other systems that need this information. These include:
 - the Service Quality Assessment (SQA) system for post-track metrics reporting
 - the complexity-based link assignment software for DSN Follow-the-Sun operations[10], which makes use of downtime, critical events, and holiday information to provide context to station operators for the links they are managing




Fig. 7. DSN This Week shows augmented level of support activities (JNO), other tracking activities for that mission, and viewperiod alignment (bottom segment). Not shown are metrics of tracking time by mission, and a table of current and upcoming downtimes by complex.

E. Forecasting and Long-Range Planning

A key function of LAPS is to generate forecast loading projections based on a variety of inputs, including mission needs and asset availabilities and capabilities. There are four kinds of alternative scenarios or studies that are typical:

1. **Changes in mission requirements:** based on new information, often from in-flight experience, there are frequently changes to mission needs that emerge and need to be assessed. Missions submit updates to their planned DSN usage which can be merged with other user inputs to look for potential contention.
2. **New missions:** usually associated with proposal calls, missions that are proposing to use the DSN submit their plans for an analysis in the context of all other expected users.
3. **Downtime studies:** DSN antennas and equipment undergo periodic planned maintenance and upgrades, sometimes requiring months of unavailability that can affect user activities. Studies of alternatives are important to help plan downtimes with the least negative user impact.
4. **Asset availability changes:** as antennas are constructed or decommissioned, the impact on mission activities needs to be assessed. This kind of study provides for comparisons of alternatives such as the use of partner agency assets to help with periods of high contention.

In LAPS, alternative scenarios can be defined from variants of the following information sets:

- assets - antennas, equipment, and their capabilities (such as what bands are supported)
- service users - missions and their requested service capabilities (such as bands required)
- downtime - projected downtime per antenna over the period of interest
- planning parameters - constraint definitions and MSPA groups

Downloaded by NASA JET PROPULSION LABORATORY on June 10, 2018 | http://arc.aiaa.org | DOI: 10.2514/6.2018-2728

All such variants are versioned, for traceability and to minimize inadvertent data loss. In addition, multiple mission scenarios (or their variants) can be incorporated into a single forecast run, each reflecting a specific set of requirements. A forecast run is typically broken into segments of at most one year duration, to keep the run time less than ~5 minutes. A forecast run goes through the following computational steps:

1. For each mission planning request, determine:
 - constraining time range
 - potential assets (single antenna) or asset sets (multi-antenna configurations, such as arrays)
 - visibility intervals from all potential assets, accounting for special elevation limits, if any
2. Accumulate with (initially) equal weight the distributed requested track duration over all the potential asset options
3. Search over all oversubscribed intervals and adjust weights to lower oversubscription on each asset without making oversubscription worse on any other asset, stopping when there are no further improvements
4. Tabulate oversubscription by asset group and by mission, for aggregation and reporting

Special processing is done for downlink-only requirements in MSPA groups, since up to four can co-exist with uplink activities without oversubscribing an antenna. Other special processing directives can be used to vary the scenario, for example:

- Introduce a “new” antenna as a hypothetical clone of an existing one, for assessing the impact of bringing additional assets to bear (such as partner agency antennas)
- Shift alternatives that allow either single antenna or arrays (e.g. see Fig. 3) towards one or the other of these alternatives, to assess the tradeoff between less time on larger aperture antennas vs. more overall antenna time but constrained to smaller aperture 34-meter antennas.
- For any requests that can be satisfied by use of *either* a 70-meter or a 34-meter antenna, adjust the usage to free time on the 70-meter assets for those missions that can *only* use the larger aperture antenna. This will reduce oversubscription on the 70-meter antennas, at the possible expense of the 34-meter antennas.

Forecast results are captured in the LAPS database where they may be viewed in graphical and tabular form or exported (Excel) at any time. All forecast runs are traceable to the precise versions of the input scenarios and run directives as described above, for repeatability and traceability in case further analysis is needed.

III. Status and Plans

LAPS v1.0 has been deployed for initial use in early 2018, offering multiple advantages over the legacy software it replaces:

- LAPS is fast (~5 minutes/forecast year) which allows more alternatives to be explored
- LAPS modeling is more accurate and reflects such capabilities as MSPA and non-standard viewperiod elevations
- All LAPS data is collected in a single unified database where it is versioned and linked to all dependent usages, including all data that can affect the results. This will enable change tracking of user requests as a gauge of system demand.
- The LAPS data model is close to that of S³, as well as being based on a large amount of shared software, thus enabling the export of planning requests from LAPS to S³
- LAPS supports direct online editing of planning requests, rather than transcription from spreadsheets; with training, individual mission users will be able to enter, update, and baseline their own planning inputs

The implementation of LAPS has taken place over several years, and some of the lessons learned over the design and implementation process include:

- Early in the development cycle it became clear that management of data *associated* with forecasting was crucial to the success of the forecasting effort itself. As a consequence, initial efforts focused on downtime planning and capturing mission critical events, and these were deployed more than a year before LAPS itself.
- In the same vein, deploying *incremental* capabilities has been important in getting users familiar with the system in steps, and in obtaining feedback along the way to incorporate in the final product.
- While LAPS represents a radical change in the DSN planning and forecasting tool suite, it has retained a modified version of the User Loading Profile (ULP) that users have become familiar with over literally decades of use. In the past, the ULP was the mission input to the long-range planning process; with LAPS it has become a report, but one that is entirely familiar. This allows users to check their expectations of the new software based on the familiar form of the ULP, and thus has helped lower the barrier to acceptance.

Next steps for LAPS include extended reporting, comparison to mid-range and executed schedules, and more comprehensive viewperiod management tools. In addition, a training exercise is planned that will allow mission users to directly log in and manage their own mission's data.

IV. Conclusions

We have described the DSN end-to-end planning and scheduling system, focusing on the newly deployed long-range planning and forecasting component known as LAPS (Loading Analysis and Planning Software). LAPS provides the capability to do extended loading studies, run what-if and trade-off analyses, and assess alternative scenarios for downtime, new missions, and new or changed assets. The LAPS long-term planning component integrates with and complements the mid- and near real-time scheduling component, Service Scheduling Software (S³), and the Service Preparation Subsystem (SPS) responsible for predict generation for execution. The delivery of LAPS completes the last stage of the plan to replace the DSN legacy planning and scheduling tool collection with a modern, integrated, web-based suite of user-driven applications.

Acknowledgments

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- [1] M. D. Johnston, D. Tran, B. Arroyo, J. Call, and M. Mercado, "Request-Driven Schedule Automation for the Deep Space Network," presented at the SpaceOps 2010, Huntsville, AL, 2010.
- [2] M. D. Johnston *et al.*, "Automating Mid- and Long-Range Scheduling for NASA's Deep Space Network," presented at the SpaceOps 2012, Stockholm, Sweden, 2012.
- [3] J. Carruth, M. D. Johnston, A. Coffman, M. Wallace, B. Arroyo, and S. Malhotra, "A Collaborative Scheduling Environment for NASA's Deep Space Network," presented at the SpaceOps 2010, Huntsville, AL, 2010.
- [4] W. A. Webb, "Forecasting of loading on the Deep Space Network for proposed future NASA mission sets," presented at the American Astronautical Society and American Institute of Aeronautics and Astronautics, Astrodynamics Specialist Conference, Provincetown, MA, 1979.
- [5] S. J. Loyola, "PC4CAST: A Tool for DSN Load Forecasting and Capacity Planning," The Telecommunications and Data Acquisition Report, p 170–184, 1993.
- [6] D. Wertz, S. Loyola, and S. Zendejas, "FASTER - A tool for DSN forecasting and scheduling," presented at the AIAA Computing in Aerospace Conference, 9th, San Diego, CA, 1993.
- [7] D. P. Holmes, "Capacity Planning for the Future Space Mission Environment," JPL, Technical Report 1997-09–02, 1997.
- [8] N. Lacey and D. G. Morris, "JPL RAPSO Long-Range Forecasting," presented at the 12th AAS/AIAA Space Flight Mechanics Meeting, San Antonio, Texas, 2002.
- [9] S. Lichten, D. Abraham, B. Arroyo, S. Asmar, J. Bell, and C. Edwards, "Allocation of Deep Space Network Ground System Tracking and Communications Assets During the 2020-2021 Timeframe of the 'Mars Armada,'" presented at the SpaceOps, Marseilles, France, 2018.
- [10] J. Lad, M. D. Johnston, D. Tran, D. Brown, K. Roffo, and C. Lee, "Complexity-Based Link Assignment for NASA's Deep Space Network for Follow-the-Sun Operations," presented at the SpaceOps, Marseilles, France, 2018.